

Targeting Optimal Active Learning via Example Quality

Lewis P. G. Evans

LEWIS.EVANS10@IMPERIAL.AC.UK

*Department of Mathematics
Imperial College London
London, SW7 2AZ, United Kingdom*

Niall M. Adams

N.ADAMS@IMPERIAL.AC.UK

*Department of Mathematics
Imperial College London
London, SW7 2AZ, United Kingdom
Heilbronn Institute for Mathematical Research
University of Bristol
PO Box 2495, Bristol, BS8 9AG, United Kingdom*

Christoforos Anagnostopoulos

CANAGNOS@IMPERIAL.AC.UK

*Department of Mathematics
Imperial College London
London, SW7 2AZ, United Kingdom*

Editor: Yoav Freund

Abstract

In many classification problems unlabelled data is abundant and a subset can be chosen for labelling. This defines the context of active learning (AL), where methods systematically select that subset, to improve a classifier by retraining.

Given a classification problem, and a classifier trained on a small number of labelled examples, consider the selection of a single further example. This example will be labelled by the oracle and then used to retrain the classifier. This example selection raises a central question: given a fully specified stochastic description of the classification problem, which example is the optimal selection?

If optimality is defined in terms of loss, this definition directly produces expected loss reduction (ELR), a central quantity whose maximum yields the optimal example selection. This work presents a new theoretical approach to AL, *example quality*, which defines optimal AL behaviour in terms of ELR.

Once optimal AL behaviour is defined mathematically, reasoning about this abstraction provides insights into AL. In a theoretical context the optimal selection is compared to existing AL methods, showing that heuristics can make sub-optimal selections.

Algorithms are constructed to estimate example quality directly. A large-scale experimental study shows these algorithms to be competitive with standard AL methods.

Keywords: active learning, example quality, expected loss reduction, classification

1. Introduction

Classification is a central task of statistical inference and machine learning. In certain cases unlabelled data is plentiful, and a small subset can be queried for labelling. Active learning (AL) seeks to intelligently select these unlabelled examples, to improve a base classifier.

Examples include medical image diagnosis and document categorisation (Dasgupta and Hsu, 2008; Hoi et al., 2006). This work focusses on the selection of a single unlabelled example, and assumes that a perfect oracle supplies the labels for selected examples.

Most AL methods are heuristic, alongside a few theoretical approaches reviewed by Settles (2009); Olsson (2009). AL method performance is often assessed by large-scale experimental studies (Guyon et al., 2011; Kumar et al., 2010; Evans et al., 2013).

A prototypical AL scenario consists of a classification problem and a classifier trained on a small labelled dataset. The classifier may be improved by retraining on a single selected example. Consider the central motivating question: given a fully specified stochastic description of the classification problem, which example is the optimal selection?

Performance in classification and AL is judged by loss functions such as those described in Hand (1997), suggesting that optimality in AL selection should be defined by loss. If optimality is defined by classifier loss, then the optimal selection is the example delivering greatest loss reduction. Expected loss reduction (ELR) formalises the loss reduction that AL provides. By construction, the example that maximises ELR delivers the greatest loss reduction, and is thereby the optimal AL selection in terms of classifier loss. In this sense ELR provides a *theoretical target* for AL: a theoretical quantity whose estimation will optimise performance.

This work presents a new theoretical approach to AL, example quality (EQ), where optimal AL behaviour is defined, then explored and estimated. EQ first considers the central motivating question, then defines optimal in terms of classifier loss, since loss defines classification performance. This definition of optimality as loss reveals ELR as the central theoretical quantity, whose maximum yields the optimal selection. Having defined optimal AL behaviour as a mathematical abstraction, this abstraction is then explored, which generates new insights into AL.

For a classification problem with a fully specified stochastic description, EQ provides an exploration of ELR and the optimal selection. The performance of an AL method is difficult to calculate analytically, since the selected example is usually a complicated function of the labelled data. Even a comprehensive experimental study may fail to elucidate the performance, since there are many sources of variation including classification problem, classifier and loss. But a full stochastic description of the problem allows examination of the AL method’s selection, and comparison to the optimal selection. This comparison is made for random selection and Shannon entropy in Section 3.2 (both are defined in Section 2.2).

EQ addresses applications by motivating the statistical estimation of ELR by new algorithms. Algorithm development reveals surprisingly complex issues which will hopefully motivate further work to improve EQ estimation.

A large-scale experimental study evaluates the performance of EQ-estimation algorithms, alongside standard AL methods from the literature. The study explores several sources of variation: multiple classifiers, AL algorithms, and real and theoretical classification problems (both binary and multi-class). The results show that the EQ-motivated algorithms perform competitively with standard AL methods. This study finds that no single AL method is a panacea, a conclusion shared by other extensive studies of AL (Guyon et al., 2011; Kumar et al., 2010; Evans et al., 2013).

The optimal AL behaviour ELR and many AL methods condition on the labelled dataset, via their dependence on the classifier. To address this dependence, EQ generalises ELR in expectation over the labelled data; this defines optimal AL behaviour for the classification problem itself, described in Sections 3.1.1 and 3.2.3. This dependence also raises the new question of the robustness of the optimal AL selection to changes in the labelled data, examined in Section 3.3.

This work is structured as follows: first the background of classification and AL are described in Section 2. Section 3 defines EQ, illustrated by a theoretical classification problem in Section 3.2. EQ estimation algorithms are described in Section 4 and evaluated in the large-scale experimental study of Section 5, followed by concluding remarks.

2. Background

The background contexts of classification and AL are described, followed by a brief review of relevant literature, with particular focus on methods that are used later in the paper.

2.1 Classification

A somewhat non-standard notation is developed to support this work, which stresses the dependence of the classifier on the training data. The categorical response variable Y is modelled as a function of the covariates \mathbf{X} . For the response Y there are k classes with class labels $\{c_1, c_2, \dots, c_k\}$. Each classification example is denoted (\mathbf{x}, y) , where \mathbf{x} is a d -dimensional covariate vector and y is a class label. The class prior is denoted $\boldsymbol{\pi}$.

The Bayes classifier is an idealisation based on the true distributions of the classes, thereby producing optimal probability estimates, and class allocations given a loss function. Given a covariate vector \mathbf{x} , the Bayes classifier outputs the class probability vector of $Y|\mathbf{x}$ denoted $\mathbf{p} = (p_j)_1^k$. A *probabilistic* classifier estimates the class probability vector as $\hat{\mathbf{p}} = (\hat{p}_j)_1^k$, and allocates \mathbf{x} to class \hat{y} using decision theoretic arguments, often using a probability threshold. This allocation function is denoted h : $\hat{y} = h(\hat{\mathbf{p}})$. For example, to minimise misclassification error, the most probable class is allocated: $\hat{y} = h(\hat{\mathbf{p}}) = \arg \max_j (\hat{p}_j)$. The objective of classification is to learn a rule with good generalisation properties.

A dataset is a set of examples, denoted $D = \{\mathbf{x}_i, y_i\}_{i=1}^n$, where i indexes the example. This indexing notation will be useful later in the exposition. A dataset D_Z may be subdivided into training data D_T and test data D_E . This dataset division may be represented by index sets, for example, given an index set $Z = \{1, \dots, n_z\}$ with (index) subsets T and E , then $Z = T \cup E$ and $D_Z = D_T \cup D_E$ show the data division into training and test subsets.

First consider a parametric classifier, for example linear discriminant analysis or logistic regression (Bishop, 2007, Chapter 4). A parametric classifier has estimated parameters $\hat{\boldsymbol{\theta}}$, which can be regarded as a fixed length vector (fixed given d and k). These parameters are estimated by model fitting to the training data: $\hat{\boldsymbol{\theta}} = \theta(D_T)$, where θ is the model fitting function. This notation is intended to emphasize the dependence of the estimated parameters $\hat{\boldsymbol{\theta}}$ on the training data D_T .

Second, this notation is slightly abused to extend to non-parametric classifiers. The complexity of non-parametric classifiers may increase with sample size, hence they cannot be represented by a fixed length object. In this case $\hat{\boldsymbol{\theta}}$ becomes a variable-length object containing the classifier’s internal data (for example the nodes of a decision tree, or the

stored examples of K -nearest-neighbours). While the contents and meaning of $\hat{\theta}$ would be very different, the classifier’s functional roles are identical: model training produces $\hat{\theta}$, which is used to predict class probabilities. These predictions are in turn used to assess classifier performance.

To consider classifier performance, first assume a fixed training dataset D_T . Classifier performance is assessed by a loss function, for example error rate, which quantifies the disagreement between the classifier’s predictions and the truth. The empirical loss for a single example is defined via a loss function $g(y, \hat{\mathbf{p}})$. Many loss functions focus on the allocated class, for example error rate $g_e(y, \hat{\mathbf{p}}) = \mathbb{1}(y \neq h(\hat{\mathbf{p}}))$. Other loss functions focus on the predicted probability, for example log loss $g_o(\hat{\mathbf{p}}) = \sum_{j=1}^k (p_j \log \hat{p}_j)$.

The estimated probabilities $\hat{\mathbf{p}}$ are highly dependent on the estimated classifier $\hat{\theta}$. To emphasize that dependence, the empirical loss for a single example is denoted $M(\hat{\theta}, \mathbf{x}, y) = g(y, \hat{\mathbf{p}})$.

To address the classifier’s expected future loss, empirical loss is generalised to expected loss, denoted $L(\hat{\theta})$:

$$L(\hat{\theta}) = E_{\mathbf{X}, Y}[M(\hat{\theta}, \mathbf{x}, y)] = E_{Y|\mathbf{X}} E_{\mathbf{X}}[M(\hat{\theta}, \mathbf{x}, y)].$$

This expected loss L is defined as an expectation over all possible test data. The expected error rate and log loss are denoted L_e and L_o . Hereafter loss will always refer to the expected loss L . The loss L is dependent on the data D used to train the classifier, emphasized by rewriting $L(\hat{\theta})$ as $L(\theta(D))$ since $\hat{\theta} = \theta(D)$.

The change in the loss as the number of labelled examples increases is of great methodological interest. This function is known as the *learning curve*, typically defined as the change of expected loss with the number of examples. Learning curves are illustrated in Figure 1, and discussed in Perlich et al. (2003); Gu et al. (2001); Kadie (1995).

Section 5 describes experiments with four classifiers: linear discriminant analysis, K -nearest-neighbours, naïve Bayes and support vector machine. Linear discriminant analysis (LDA) is a linear generative classifier described in Hastie et al. (2009, Chapter 4). K -Nearest-Neighbours (K -nn) is a well-known non-parametric classifier discussed in Duda et al. (2001, Chapter 4). Naïve Bayes is a probabilistic classifier which assumes independence of the covariates, given the class; see Hand and Yu (2001). The support vector machine (SVM) is a popular non-parametric classifier described in Cortes and Vapnik (1995). Standard R implementations are used for these classifiers.

2.2 Active Learning

The context for AL is a scarcity of labelled data but an abundance of unlabelled examples. Good introductions to AL are provided by Dasgupta (2011), Settles (2009) and Olsson (2009).

An algorithm can select a few unlabelled examples to obtain their labels from an oracle (for example a human expert). This provides more labelled data which can be included in the training data, potentially improving a classifier. Intuitively some examples may be more informative than others, so systematic example selection should maximise classifier improvement.

In *pool-based* AL, there is an unlabelled pool of data X_P from which examples may be selected for labelling. This pool provides a set of examples for label querying, and also gives further information on the distribution of the covariates. Usually there is also a (relatively small) dataset of labelled examples, denoted D_S , of size n_s . This work considers the scenario of pool-based AL.

In AL it is common to examine the learning curve, by repeating the AL selection step many times (*iterated AL*). At each selection step, the loss is recorded, and this generates a set of losses, which define the learning curve for the AL method. Iterated AL allows the exploration of performance over the learning curve, as the amount of labelled data grows. This repeated application of AL selection is common in both applications and experimental studies (Guyon et al., 2011; Evans et al., 2013). A recurring theme in AL is the strong dependence of the classifier on the training data, and iterated AL provides a detailed examination of this dependence.

At each selection step, an AL method may select a single example from the pool (*sequential AL*) or several examples at once (*batch AL*). AL applications are often constrained to use batch AL for pragmatic reasons (Settles, 2009).

Turning to AL performance, consider random selection (RS) where examples are chosen randomly (uniformly) from the pool. By contrast, AL methods select some examples in preference to others. Under RS and AL, the classifier receives exactly the same number of labelled examples; thus RS provides a reasonable benchmark for AL (Guyon et al., 2011; Evans et al., 2013). The comparison of methods to benchmarks is available in experiments but not in AL applications (Provost and Attenberg, 2010).

Classifier performance (usually) improves even under the benchmark RS, since the classifier receives more training data. AL performance assessment should consider how much AL outperforms RS. Hence AL performance addresses the relative improvement of AL over RS, and the relative ranks of AL methods, rather than the absolute level of classifier performance. Figure 1 shows the losses of AL and RS as the number of labelled examples increases.

Figure 1 shows two different senses in which AL outperforms RS: first AL achieves better loss reduction for the same number of labels (*fixed-label comparison*), and second AL needs fewer labels to reach the same classifier performance (*fixed-loss comparison*). Together the fixed-label comparison and fixed-loss comparison form the two fundamental aspects of AL performance. The fixed-label comparison first fixes the number of labels, then seeks to minimise loss. Several established performance metrics focus on the fixed-label comparison: AUA, ALC and WI (Guyon et al., 2011; Evans et al., 2013). The fixed-label comparison is more common in applications where the costs of labelling are significant (Settles, 2009).

Under the fixed-loss comparison, the desired level of classifier loss is fixed, the goal being to minimise the number of labels needed to reach that level. Label complexity is the classic example, where the desired loss level is a fixed ratio of asymptotic classifier performance (Dasgupta, 2011). Label complexity is often used as a performance metric in contexts where certain assumptions permit analytically tractable results, for example Dasgupta (2011).




Figure 1: Performance comparison of active learning and random selection, showing that a classifier often improves faster under AL than under RS. In both cases the loss decreases as the number of labelled examples increases; however, AL improves faster than RS. These curves are smoothed averages from multiple experiments. The black vertical line illustrates the fixed-label comparison, whereas the blue horizontal line shows the fixed-loss comparison (see Section 2.2). The classification problem is “Abalone” from UCI, a three-class problem, with the base classifier being 5-nn, and the AL method being Shannon entropy.

2.3 Overview of Active Learning Methods

Uncertainty sampling is a heuristic approach to AL, where examples are chosen closest to the classifier’s estimated decision boundary (Thrun and Möller, 1992; Settles, 2009). This approach selects examples of the greatest classifier uncertainty in terms of class membership probability. The idea is that these uncertain examples will be the most useful for tuning the classifier’s decision boundary. Example methods include Shannon entropy (SE), least confidence and maximum uncertainty. For a single unlabelled example \mathbf{x} , least confidence is defined as $U_L(\mathbf{x}, \theta(D)) = 1 - \hat{p}(\hat{y}|\mathbf{x})$, where $\hat{p}(\hat{y}|\mathbf{x})$ is the classifier’s estimated probability of the allocated class \hat{y} . Shannon entropy is defined as $U_E(\mathbf{x}, \theta(D)) = \sum_{j=1}^k \hat{p}_j \log(\hat{p}_j)$. The uncertainty sampling approach is popular and efficient, but lacks theoretical justification.

Version space search is a theoretical approach to AL, where the version space is the set of hypotheses (classifiers) that are consistent with the data (Mitchell, 1997; Dasgupta, 2011). Learning is then interpreted as a search through version space for the optimal hypothesis. The central idea is that AL can search this version space more efficiently than RS.

Query by committee (QBC) is a loose approximation to version space search (Seung et al., 1992). Here a committee of classifiers is trained on the labelled data, which then selects the unlabelled examples where the committee’s predictions disagree the most. This prediction disagreement may focus on predicted classes (for example vote entropy) or predicted class probabilities (for example average Kullback-Liebler divergence); see Olsson (2009). These widely used versions of QBC are denoted QbcV and QbcA. A critical control parameter for QBC is the choice of classifier committee, which lacks theoretical guidance. In this sense version space search leaves the optimal AL selection unspecified.

Another approach to AL is exploitation of cluster structure in the pool. Elucidating the cluster structure of the pool could provide valuable insights for example selection. Dasgupta (2011) gives a motivating example: if the pool clusters neatly into b class-pure clusters where $b = k$, then b labels could suffice to build an optimal classifier. This very optimistic example does illustrate the potential gain.

A third theoretical approach is error reduction, introduced in Roy and McCallum (2001). This approach minimises the error of the retrained classifier, which is the error of the classifier which has been retrained on the selected example. Those authors consider two loss functions, error rate and log loss, to construct two quantities, which are referred to here as expected future error (EFE) and expected future log loss (EFL). Those authors focus on methods to estimate EFE and EFL, before examining the experimental performance of their estimators.

The error reduction approach is similar in spirit to example quality, since the optimal example selection is first considered, and then specified in terms of classifier loss. The quantity EFE is a valuable precursor to expected loss reduction, being a component of ELR, which is defined in Equation 2. However EFE omits the loss of the current classifier, which proves important when taking expectation over the labelled data (see Sections 3 and 3.3) and when examining improvement (see Section 3.2). Those authors focus on EFE estimation, rather than using EFE and EFL to construct a theoretical account of AL.

Given a classifier fitting function θ , labelled data D and a single unlabelled example \mathbf{x} , EFE is defined as

$$EFE(\mathbf{x}, \theta, D) = -E_{Y|\mathbf{x}}[L_e(\theta(D \cup (\mathbf{x}, Y)))] = -\sum_{j=1}^k \{p_j L_e(\theta(D \cup (\mathbf{x}, c_j)))\},$$

where L_e is error rate (see Section 2.1). EFL is defined in the very same way as EFE, with log loss L_o replacing error rate L_e . Both of these quantities average over the unobserved label $Y|\mathbf{x}$.

Those authors define an algorithm to calculate EFE, denoted EfeLc, which approximates the loss using the unlabelled pool for efficiency. Specifically it approximates error rate L_e by the total least confidence over the entire pool:

$$L_e(\theta(D)) \approx \sum_{\mathbf{x}_i \in X_P} U_L(\mathbf{x}_i, \theta(D)),$$

where X_P are the unlabelled examples in the pool.

Those authors propose the following approximation for the value of EFE by calculating

$$f_1(\mathbf{x}, \theta, D) = -\sum_{j=1}^k \left\{ \hat{p}_j \sum_{\mathbf{x}_i \in X_P} U_L(\mathbf{x}_i, \theta(D \cup (\mathbf{x}_i, c_j))) \right\} = -\sum_{j=1}^k \left\{ \hat{p}_j \sum_{\mathbf{x}_i \in X_P} (1 - \hat{p}(\hat{y}_i|\mathbf{x}_i)) \right\}. \quad (1)$$

This approximation of L_e by the total least confidence over the pool is potentially problematic. It is easy to construct cases (for example an extreme outlier) where a labelled example would reduce a classifier's uncertainty, but also increase the overall error; such examples call into question the approximation of error by uncertainty. In the absence of further assumptions or motivation, it is hard to anticipate the statistical properties of f_1 in Equation 1 as an estimator.

3. Example Quality

Here the theoretical target, example quality, is defined as expected loss reduction. This motivates EQ as an estimation target, both theoretically and for applications.

3.1 The Definition of Example Quality

To define the theoretical target, knowledge is assumed of the underlying distribution (\mathbf{X}, Y) , with all expectations being formed with respect to that joint distribution. Assume a fixed dataset D_S sampled i.i.d. from the joint distribution (\mathbf{X}, Y) . The dependence of the classifier $\hat{\theta}$ on the data D_S is critical, with the notation $\hat{\theta} = \theta(D_S)$ intended to emphasize this dependence.

First assume a base classifier already trained on a dataset D_S . Consider how much a single labelled example improves performance. The single labelled example (\mathbf{x}, y) will be chosen from a *labelled* dataset D_W , to develop the argument. This loss from retraining on that single labelled example is examined in order to later define the loss for the expected label of an unlabelled example.

Consider the selection of a single labelled example (\mathbf{x}, y) from D_W , given the labelled data D_S , the classifier training function θ and a loss function L . The reduction of the loss for retraining on that example is defined as actual-EQ, denoted Q^a :

$$Q^a(\mathbf{x}, y, \theta, D_S) = L(\theta(D_S)) - L(\theta(D_S \cup (\mathbf{x}, y))).$$

Q^a is the actual classifier improvement from retraining on the labelled example (\mathbf{x}, y) . The goal here is to maximise the reduction of loss. The greatest loss reduction is achieved by selecting the example (\mathbf{x}_*, y_*) from D_W that maximises Q^a , given by

$$(\mathbf{x}_*, y_*) = \arg \max_{(\mathbf{x}, y) \in D_W} Q^a(\mathbf{x}, y, \theta, D_S).$$

Turning to AL, the single example \mathbf{x} is unlabelled, and will be chosen from the unlabelled pool X_P . Knowledge of the underlying joint distribution (\mathbf{X}, Y) is still assumed. Here the unknown label of \mathbf{x} is a random variable, $Y|\mathbf{x}$, and taking its expectation allows the expected loss to be defined, this being the classifier loss after retraining with the unlabelled example and its unknown label. Thus the expected loss is defined using the expectation over the label $Y|\mathbf{x}$ to form conditional-EQ, denoted Q^c :

$$\begin{aligned} Q^c(\mathbf{x}, \theta, D_S) &= E_{Y|\mathbf{x}}[Q^a(\mathbf{x}, Y, \theta, D_S)] = L(\theta(D_S)) - E_{Y|\mathbf{x}}[L(\theta(D_S \cup (\mathbf{x}, Y)))] \\ &= \underbrace{L(\theta(D_S))}_{\text{Term F}} - \underbrace{\sum_{j=1}^k \{p_j L(\theta(D_S \cup (\mathbf{x}, c_j)))\}}_{\substack{\text{Term H} \\ \text{Term J}}}. \end{aligned} \quad (2)$$

Q^c defines the expected loss reduction from retraining on the unlabelled example \mathbf{x} with its unknown label $Y|\mathbf{x}$. In other words Q^c captures the difference between two losses, of the existing classifier against the expected loss of the retrained classifier. In this sense Q^c is the improvement function, since it defines exactly how much this example will improve the classifier.

The unlabelled example \mathbf{x}_* from the pool X_P that maximises Q^c is the optimal example selection:

$$\mathbf{x}_* = \arg \max_{\mathbf{x} \in X_P} Q^c(\mathbf{x}, \theta, D_S).$$

Algorithms to estimate the target Q^c are presented in Section 4.

The target Q^c extends to define the optimal selection for batch AL, but this extension is omitted for space.

For a theoretical classification problem, the target Q^c can be evaluated exactly, to reveal the best and worst possible loss reduction (in expectation), by maximising and minimising Q^c . Figure 2 shows that the best and worst AL performance curves are obtained by maximising and minimising Q^c .

Given a theoretical classification problem, the target Q^c also provides straightforward calculations for the performance of random selection, and the regret of AL methods such as Shannon entropy; these calculations are omitted for space.




Figure 2: The best and worst AL performance curves are obtained by maximising and minimising the target Q^c , which demonstrate the extremes of AL performance. With simulated data, Q^c can be calculated exactly; here the classification problem is the Ripley four-Gaussian problem (illustrated in Figure 6e). These curves are smoothed from multiple experiments, with the base classifier being 5-nn.

3.1.1 OPTIMAL SELECTION INDEPENDENT OF THE LABELLED DATA

Here the optimal AL behaviour is examined, independently of the labelled data D_S . This provides further insight into the optimal AL selection for the classification problem itself.

The target Q^c conditions on the data D_S , which is a primary source of variation for Q^c . To address this source of variation, the expectation over D_S is taken, leading to marginal-EQ, denoted Q^m :

$$Q^m(\mathbf{x}, \theta, n_s) = E_D[Q^c(\mathbf{x}, \theta, D)] = E_D[L(\theta(D)) - E_{Y|\mathbf{x}}[L(\theta(D \cup (\mathbf{x}, Y)))]], \quad (3)$$

where n_s is the constant size of the dataset D_S : $n_s = |D_S|$.

Q^m is the marginal improvement expected for a classifier retrained on the unlabelled example \mathbf{x} , in expectation over a dataset D_S drawn i.i.d. from (\mathbf{X}, Y) . The dataset size n_s still matters: as $n_s \rightarrow \infty$, the ranking signal of Q^m vanishes, as all examples provide equal (zero) loss reduction.

The unlabelled example that maximises Q^m is denoted \mathbf{x}_+ :

$$\mathbf{x}_+ = \arg \max_{\mathbf{x} \in X_P} Q^m(\mathbf{x}, \theta, n_s).$$

This marginal target Q^m defines optimal AL behaviour for the underlying classification problem itself, independently of D_S . The targets Q^c and Q^m , with their maxima x_* and x_+ , reveal optimal AL behaviour, whether conditioning on a single dataset (Q^c) or whether considering the classification problem as a whole (Q^m).

Q^m is calculated and illustrated for a theoretical classification problem in Section 3.2.3. The target Q^m raises another question for AL: exactly how Q^c depends on D_S , and the consequences for optimal AL behaviour; this is addressed in Section 3.3.

The central motivating question from The Introduction, which example is the optimal selection, can now be generalised to form a second central question: which example is the optimal selection, for the classification problem itself, independently of D_S ? This second question is answered by the marginal target Q^m and its maximum x_+ , illustrated in Figures 5, 6f and 9. This second question creates a new shift in AL from conditional-AL (Q^c , x_*) to marginal AL (Q^m , x_+).

3.2 Theoretical Example

An example using a theoretical classification problem is presented, to illustrate EQ in detail. The stochastic character of this problem is fully specified, allowing exact calculations of the loss L , and the theoretical targets Q^c and Q^m as functions of the univariate covariate x . To reason about Q^c as a function of x , an infinite pool is assumed, allowing any $x \in \mathbb{R}$ to be selected. These targets are then explored as functions of x , and the optimal AL selection x_* is examined, where $x_* = \arg \max_x Q^c(x)$, $x \in \mathbb{R}$.

When the full stochastic description of the problem is known, this knowledge allows examination of the AL method's selection, denoted x_r , and comparison to the optimal selection x_* . This comparison is made below for the popular AL heuristic Shannon entropy, and for random selection.

Consider a binary univariate problem, defined by a balanced mixture of two Gaussians: $\{\pi = (\frac{1}{2}, \frac{1}{2}), (X|Y = c_1) \sim N(-1, 1), (X|Y = c_2) \sim N(1, 1)\}$. The prior π is fixed here

(though varied later to illustrate the target Q^m , see Section 3.2.3). The loss function is error rate L_e (defined in Section 2.1). The true decision boundary to minimise error rate is denoted $t = \frac{1}{2}(\mu_1 + \mu_2)$.

Every dataset D of size n sampled from this problem is assumed to split equally into two class-pure subsets $D_j = \{y_i = c_j, (x_i, y_i) \in D\}$ each of size $n_j = \frac{n}{2}$; this is sampling while holding the prior fixed. Appendix A provides full calculation details for this example.

Consider a classifier that estimates only the class-conditional means, given the true prior π and the true common variance of 1. The classifier parameter vector is $\hat{\theta} = (\hat{\mu}_1, \hat{\mu}_2)$, where $\hat{\mu}_j$ is the sample mean for class c_j . The classifier's estimated decision boundary to minimise error rate is denoted $\hat{t} = \frac{1}{2}(\hat{\mu}_1 + \hat{\mu}_2)$.

3.2.1 CALCULATION AND EXPLORATION OF Q^c

Q^c is calculated exactly as a function of x , to explore Q^c as a function of x , and to examine the optimal selection x_* .

The classifier's decision rule $r_1(x)$ minimises the loss $L_e(\hat{\theta})$, and is given in terms of a threshold on the estimated class probabilities by

$$r_1(x) = \begin{cases} \hat{y} = c_1 & : \hat{p}_1(x) > \frac{1}{2}, \\ \hat{y} = c_2 & : \hat{p}_1(x) < \frac{1}{2}, \end{cases}$$

or equivalently, in terms of a decision boundary on x , by

$$r_1(x) = \begin{cases} \hat{\mu}_1 < \hat{\mu}_2 & : \hat{y} = c_1 \text{ if } x < \hat{t}, c_2 \text{ otherwise,} \\ \hat{\mu}_1 > \hat{\mu}_2 & : \hat{y} = c_1 \text{ if } x > \hat{t}, c_2 \text{ otherwise.} \end{cases}$$

The classifier may get the estimated class means the wrong way around, in the unlikely case that $\hat{\mu}_1 > \hat{\mu}_2$. As a result the classifier's behaviour is very sensitive to the condition $(\hat{\mu}_1 > \hat{\mu}_2)$, as shown by the second form of the decision rule $r_1(x)$, and by the loss function in Equation 4.

It is straightforward to show that the loss $L_e(\hat{\theta})$ is given by

$$L_e(\hat{\theta}) = \frac{1}{2} \{1 - F_1(\hat{t}) + F_2(\hat{t}) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}_2)[2F_1(\hat{t}) - 2F_2(\hat{t})]\}, \quad (4)$$

where $F_j(x)$ denotes the cdf for class-conditional $(X|Y = c_j)$. This result is derived in Appendix A.

In AL an unlabelled point x is chosen for the oracle to label, before retraining the classifier. Retraining the classifier with a single new example (x, c_j) yields a new parameter estimate denoted $\hat{\theta}'_j$, where the mean estimate for class c_j has a new value denoted $\hat{\mu}'_j$, with a new estimated boundary denoted \hat{t}'_j .

Here $\hat{\mu}'_j = (1 - z)\hat{\mu}_j + zx$ where $z = \frac{2}{n+2}$, z being an updating constant which reflects the impact of the new example on the mean estimate $\hat{\mu}_j$.

To calculate Q^c under error loss L_e , observe that the Term J from Equation 2 is $[p_1 L_e(\hat{\theta}'_1) + p_2 L_e(\hat{\theta}'_2)]$. Term F in Equation 2 is directly given by Equation 4. From Equ-

tions 2 and 4, $Q^c(x, \theta, D) = Q^c(x, \hat{\theta}) = L_e(\hat{\theta}) - [p_1 L_e(\hat{\theta}'_1) + p_2 L_e(\hat{\theta}'_2)]$, hence

$$\begin{aligned} Q^c(x, \hat{\theta}) = Q^c(x, (\hat{\mu}_1, \hat{\mu}_2)) &= \frac{1}{2} \{1 - F_1(\hat{t}) + F_2(\hat{t}) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}_2)[2F_1(\hat{t}) - 2F_2(\hat{t})]\} \\ &\quad - \frac{p_1}{2} \{1 - F_1(\hat{t}'_1) + F_2(\hat{t}'_1) + \mathbb{1}(\hat{\mu}'_1 > \hat{\mu}_2)[2F_1(\hat{t}'_1) - 2F_2(\hat{t}'_1)]\} \\ &\quad - \frac{p_2}{2} \{1 - F_1(\hat{t}'_2) + F_2(\hat{t}'_2) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}'_2)[2F_1(\hat{t}'_2) - 2F_2(\hat{t}'_2)]\}, \end{aligned}$$

where p_j , $\hat{\mu}'_j$, and \hat{t}'_j are functions of x .

Even for this simple univariate problem, $Q^c(x, \hat{\theta})$ is a complicated non-linear function of x . Given the difficulty of analytically analysing and optimising Q^c , specific cases of $\hat{\theta}$ allow exploration of Q^c , shown in Figure 3. In each specific case of $\hat{\theta}$, x_* yields greatest correction to $\hat{\theta}$ in terms of moving the estimated boundary \hat{t} closer to the true boundary t . This is intuitively reasonable since error rate is a function of \hat{t} and minimised for $\hat{t} = t$.

In the first two cases (Figures 3a and 3b), x_* is negative, to improve the classifier by reducing the overestimate of \hat{t} . In the third case (Figure 3c), $\hat{t} = t$ and here the classifier's loss L_e cannot be reduced, shown by $Q^c(x) < 0$ for all x . The fourth case (Figure 3d) is interesting because the signs of the estimated means are reversed compared to the true means, and here the most non-central x offer greatest classifier improvement. Together these cases show that even for this toy example, the improvement function Q^c is complicated and highly dependent on the estimated parameters.

The target Q^m is calculated analytically in Appendix A, and illustrated in Section 3.2.3.

3.2.2 EXPLORATION OF SHANNON ENTROPY AND RANDOM SELECTION

In this theoretical example, a fully specified stochastic description of the classification problem is available, allowing the selection made by an AL method, x_r , to be examined and compared to the optimal selection x_* . Here this comparison is made for two example selection methods, SE and RS.

SE always selects x_r at the estimated boundary \hat{t} . RS selects uniformly from the pool, assumed to be i.i.d. in AL, hence the RS selection probability is given by the marginal density $p(x)$. In contrast to Q^c and SE, RS is a stochastic selection method, with expected selection $x_r = 0$ in this problem. Figure 4 illustrates Q^c , SE and $p(x)$ as contrasting functions of x , with very different maxima.

Q^c is asymmetric in the first two cases, and symmetric for the final two. By contrast, SE and RS are always symmetric (for all possible values of $\hat{\theta}$).

In the first two cases (Figures 4a and 3b), SE selects a central x_r , thereby missing the optimal selection x_* . In the second case (Figure 3b), SE selects x_r with $Q^c(x_r) < 0$, failing to improve the classifier, whereas the optimal selection x_* does improve the classifier since $Q^c(x_*) > 0$. The third case is unusual, since $\hat{t} = t$ and this classifier's loss L_e cannot be improved, hence $Q^c(x) < 0$ for all x . In the fourth case (Figure 4d) SE makes the worst possible choice of x . In all four cases, SE never chooses the optimal point; SE may improve the classifier, but never yields the greatest improvement. These specific cases of $\hat{\theta}$ show that SE often makes a suboptimal choice for x_r , for this theoretical example.

Turning to consider RS, for these four cases of $\hat{\theta}$, the expected RS selection is a suboptimal choice of x_* for this problem. It is notable that the expected RS selection is usually

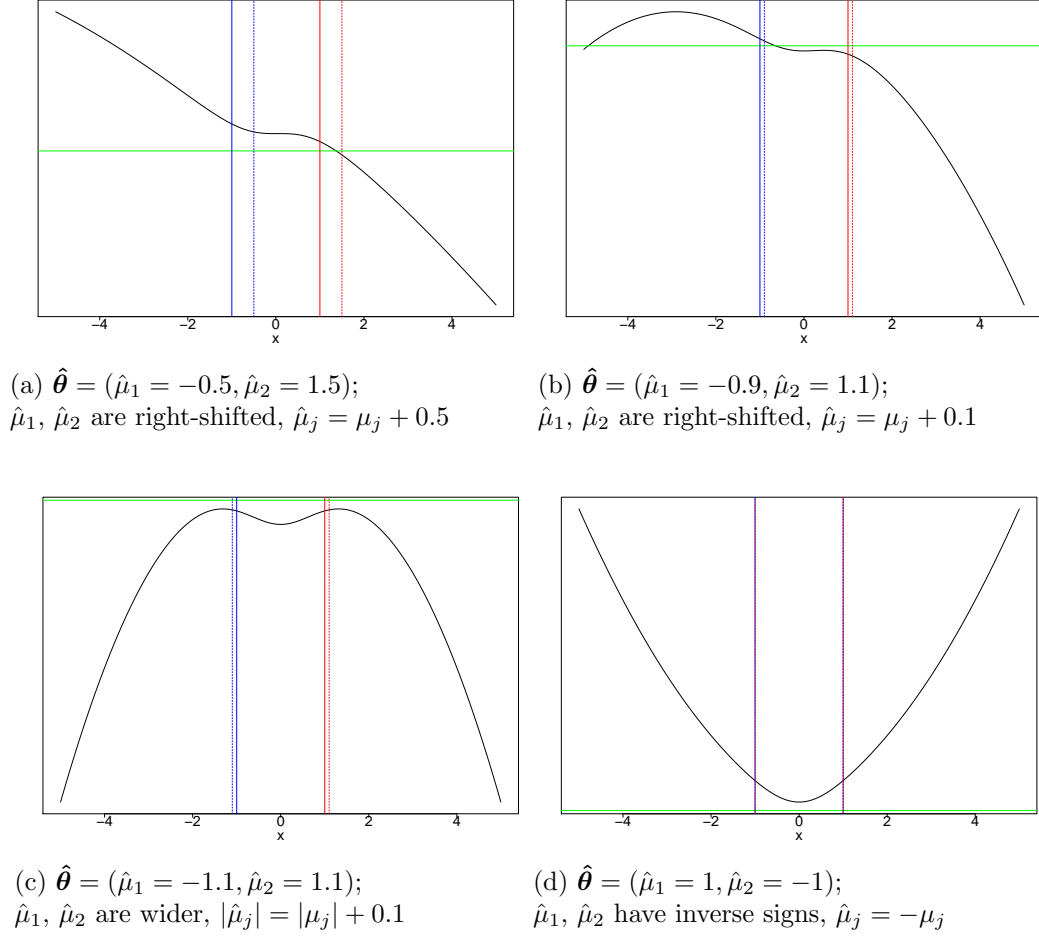


Figure 3: Illustration of the target Q^c as a function of x , for specific cases of the estimated classifier parameters $\hat{\theta} = (\hat{\mu}_1, \hat{\mu}_2)$. The class mean parameters are shown in solid blue and red, with the estimated means shown in dotted blue and red. The green line indicates $Q^c(x) = 0$ (zero improvement), with n_s being 18 in all cases. In each specific case, the optimal selection x_* yields greatest correction to $\hat{\theta}$ in terms of moving the estimated boundary \hat{t} closer to the true boundary t .




Figure 4: Comparison of Q^c against SE and RS as functions of x , for specific cases of the estimated classifier parameters $\hat{\theta} = (\hat{\mu}_1, \hat{\mu}_2)$. Q^c is shown in black, SE in purple and RS in orange (for RS, the density $p(x)$ is shown). The class mean parameters are shown in solid blue and red, with the estimated means shown in dotted blue and red. The green line indicates $Q^c(x) = 0$ (zero improvement), with n_s being 18 in all cases. The three functions are scaled to permit this comparison.

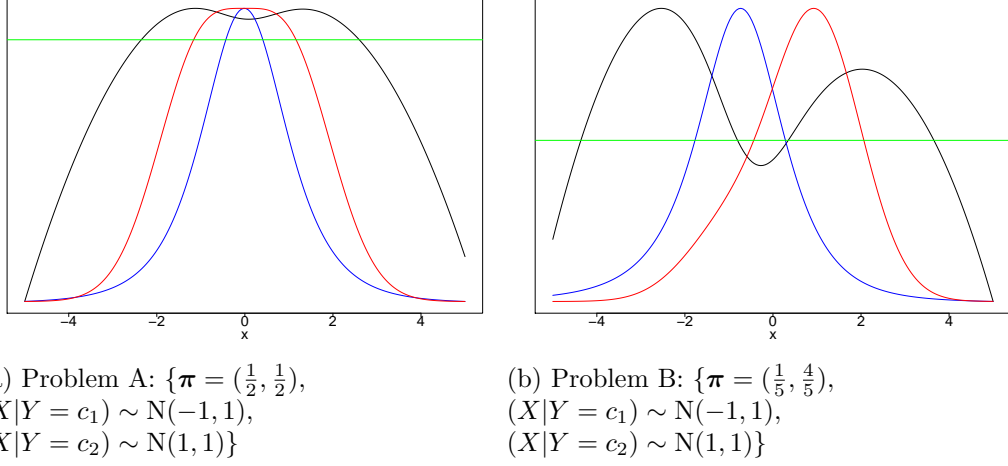


Figure 5: Comparison of SE and RS against Q^m as functions of x , with two closely related theoretical problems, which differ only in the class prior π . For both problems there are large ranges of the covariate x which improve the classifier, in expectation over D_S . Q^m is shown in black, SE in blue, and RS in red (for RS, the density $p(x)$ is shown). The green line indicates $Q^m(x) = 0$ (zero improvement), with n_s being 18 in both cases. The three functions are scaled to permit this comparison.

close to the SE selection. The stochastic nature of RS implies that it often selects far more non-central x values than SE.

3.2.3 THE AVERAGE PERFORMANCE OF SHANNON ENTROPY AND RANDOM SELECTION

Having explored specific cases of $\hat{\theta}$, the performance of AL methods on average over the labelled data D_S is now considered, by the comparison of SE and RS to the target Q^m . Q^m illuminates optimal AL selection for the classification problem itself, independently of any specific dataset D_S .

This comparison is made with two theoretical problems, the first being defined at the start of Section 3.2, and the second problem being a minor variant where the prior π has been modified to $(\frac{1}{5}, \frac{4}{5})$. Both problems are defined in Figure 5. The values for Q^m and SE are calculated by averaging over multiple i.i.d. draws of the labelled data D_S .

Figure 5 illustrates the comparison of SE and RS against Q^m with these two theoretical problems. This figure suggests that for both problems, there are large ranges of the covariate x which improve the classifier, in expectation over D_S . The locations of the improvement x values are quite different for the two problems.

3.3 The Dependence of Active Learning on the Labelled Data

Most if not all AL methods implicitly condition on the labelled data D_S , via dependence on the classifier $\hat{\theta}$, for example uncertainty sampling and QBC (described in Section 2.3). The improvement function Q^c and its maximum x_* also condition on D_S , see Equation 2. This conditioning on D_S is very natural in AL, yet is rarely made explicit, or examined.

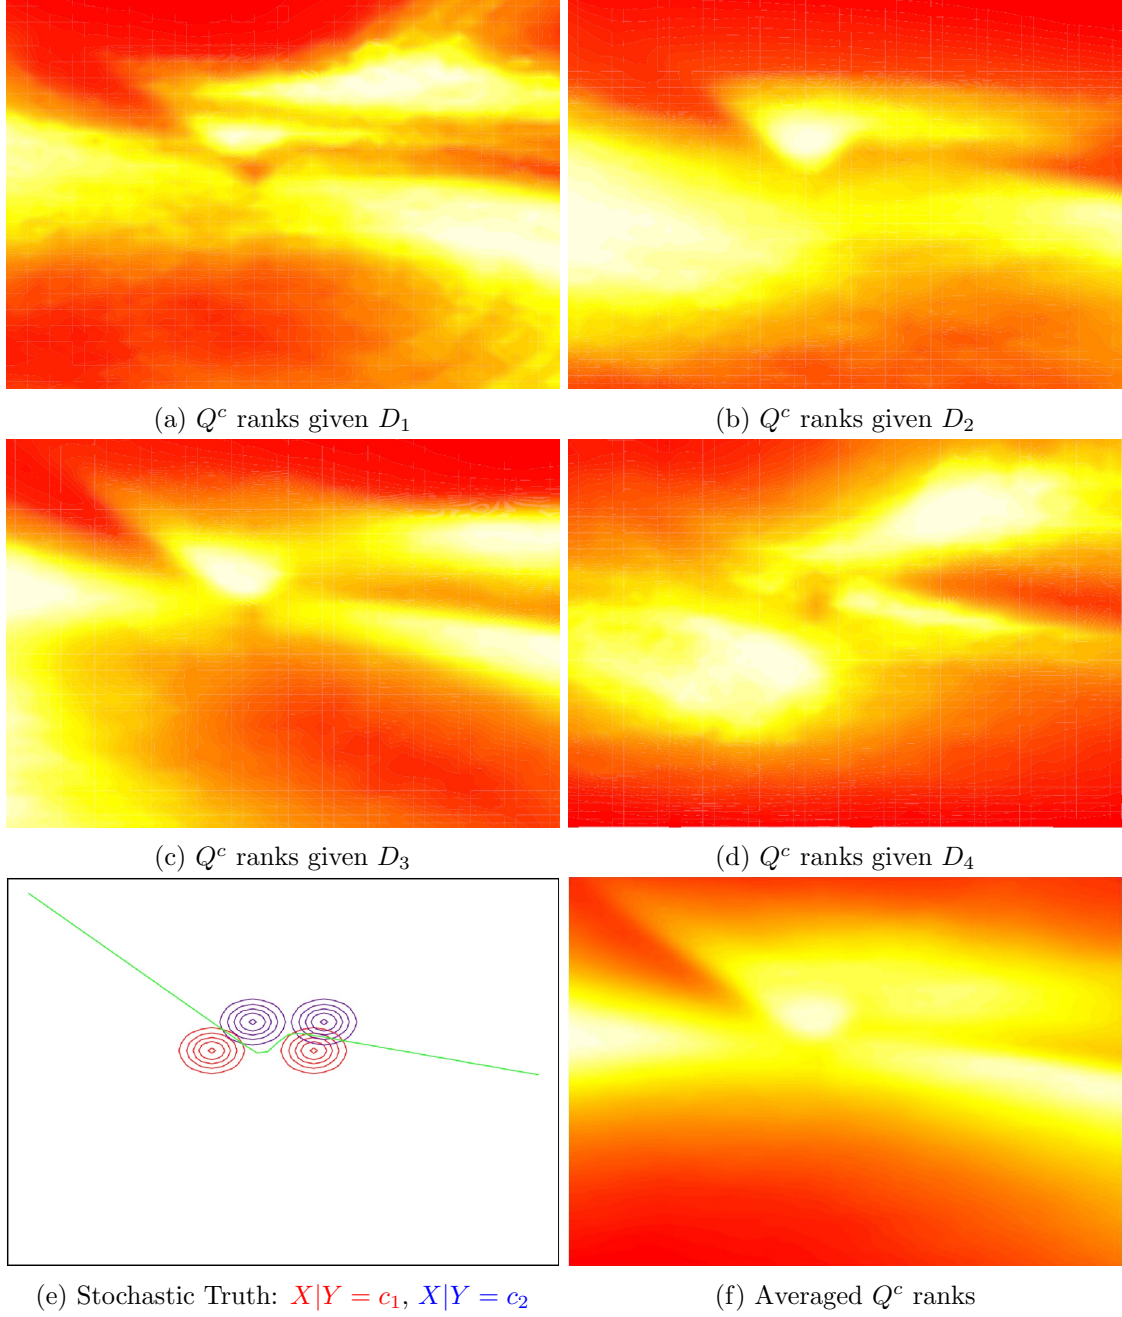


Figure 6: The visual similarity of the four Q^c rankings (Figures 6a to 6d) show the low dependence of AL on the labelled data, in other words the low sensitivity of Q^c to D_S . Ranking of Q^c values for different draws of the fixed-size dataset D_i are shown in Figures 6a to 6d. Higher ranks are shown in brighter yellow, lower ranks are darker red. The classification problem is the Ripley four-Gaussian problem, where the stochastic truth is shown in Figure 6e, see Appendix B. The size of D_i is always 20.

This dependency on D_S first motivates the definition of Q^m in Equation 3, which addresses the dependence taking the expectation over D_S . However the dependence on D_S also raises another statistical question: how does optimal AL selection depend on the labelled data D_S ?

To illustrate, one possibility is that Q^c and x_* vary wildly with D_S ; in this case (denoted *fragile*), different draws of D_S would lead to very different x_* values. Alternatively x_* might have low sensitivity to D_S , and in this second case (denoted *robust*), different draws of D_S would produce similar x_* values. Figure 6 shows Q^c and x_* for different draws of D_S , where the similarity of Figures 6a to 6d and 6f show the low sensitivity of Q^c to D_S .

Each draw of D_S provides a classifier $\hat{\theta}$, which in turn provides a specific Q^c ranking of all the examples in the pool X_P , shown in Figure 6. In the second robust case these pool rankings will be similar to each other, and this is observed visually in Figure 6. To examine this statistically, experiments examine the sensitivity of Q^c to D_S . The similarity of the pool rankings from different draws of D_S express the sensitivity of Q^c to D_S , where greater similarity shows lower sensitivity.

Here a different theoretical classification problem is considered, the Ripley four-Gaussian problem (shown in Figure 6e). The classifier is quadratic discriminant analysis (Hastie et al., 2009, Chapter 4), with the loss function being error rate L_e .

A single constant pool of unlabelled examples X_P is chosen as a fixed grid of points in covariate space, with a 2-d grid providing visualisation in Figure 6. Multiple datasets $(D_i)_1^n$ are drawn i.i.d. from (\mathbf{X}, Y) . Each dataset D_i produces a set of Q^c values for the grid, denoted \mathbf{q}_i , implying a ranking of the pool, denoted \mathbf{r}_i , shown in Figure 6.

The similarity of the pool rankings $(\mathbf{r}_i)_1^n$ is examined by standard statistical tests of correlations (with Holm and Bonferroni corrections), Moran’s I and Geary’s C (Moran, 1950; Geary, 1954). These statistical tests show that the ranks are very closely related for different draws of $(D_i)_1^n$.

This similarity of the ranks shows that, in this toy example, Q^c has low dependence on the specific dataset D_S , which is the robust case. This in turn suggests that the optimal AL choice x_* for one dataset D_i is near-optimal for a different dataset D_j .

4. Algorithms to Estimate Example Quality

Section 3 defines the targets Q^c and Q^m in a theoretical context with a full stochastic description of the classification problem. AL applications present a very different context, where all theoretical quantities must be estimated from the single labelled dataset D_S . In AL applications, there is only one label budget, and as a result AL methods cannot be compared to each other, or to benchmarks such as RS (Provost and Attenberg, 2010). The estimation of multiple theoretical quantities from the single dataset D_S introduces statistical issues described below.

The definition of Q^c in Equation 2 includes three primary components: \mathbf{p} , $\hat{\theta}$ and L . Q^c estimation requires estimating these three components from one labelled dataset D_S . Estimating three quantities from a single dataset raises interesting statistical choices. One major choice must be made between using the same data for all three estimations, termed *naïve reuse*, or partitioning the data into subsets for each estimation. This choice between

naïve reuse and partitioning has implications for the bias and variance of \hat{Q}^c estimates, discussed below.

Notationally three datasets are denoted D_C , D_T and D_E to estimate three components of Q^c :

- The class probability vector, $\mathbf{p} = p(Y|\mathbf{x})$, estimated by $\hat{\mathbf{p}}$ using data D_C ,
- The classifier parameters, $\hat{\boldsymbol{\theta}} = \theta(D_S)$, estimated by $\theta(D_T)$ using data D_T ,
- The loss, L , estimated by \hat{L} using data D_E .

Each of these three datasets must be subsets of D_S .

In the case of naïve reuse, all three datasets equal D_S . For data partitioning, the three datasets are disjoint partitions of D_S . Both of these cases are shown in Figure 7. More sophisticated partitioning schemes are under investigation and the subject of future work.

4.1 Two EQ Estimation Algorithms

Two algorithms are presented that estimate Q^c directly. The first algorithm takes a simple approach where all of D_S is used to estimate all three components. The intention is to reduce the variance of the component estimates, by using the maximum amount of data for each estimate. This first algorithm is termed simpleEQ; its naïve reuse of the data is illustrated in Figure 7. Here $D_S = D_C = D_T = D_E$.

For practical estimation of Q^c , Term F in Equation 2 can be ignored since it is independent of \mathbf{x} . Thus the main task of practical Q^c estimation is estimation of Term J in Equation 2, Term J being the expected classifier loss after retraining on the new example \mathbf{x} with its unknown label $Y|\mathbf{x}$.

The simpleEQ algorithm immediately encounters a problem in estimating Term J: the same data D_S is used both to train the classifier and also to estimate the loss. This in-sample loss estimation is known to produce optimistic, biased estimates of the loss (Hastie et al., 2009, Chapter 7).

This problem of biased component estimates under naïve reuse motivates the development of a second algorithm, termed partitionEQ. The intention is to produce unbiased component estimates, and thereby reduce the bias of the \hat{Q}^c estimate. This algorithm randomly partitions the data D_S into three disjoint subsets D_C , D_T and D_E ; this partitioning is shown in Figure 7.

Each subset is used for a single estimation task: to estimate \mathbf{p} , the classifier parameters $\hat{\boldsymbol{\theta}}$, and the loss L respectively. The random partitioning of D_S into subsets is arbitrary. For this reason, the random partitioning is performed several times, and the resulting \hat{Q}^c estimates are averaged. In the experimental study of Section 5, the partitioning is repeated ten-fold.

Random sub-sampling of the pool is used for computational efficiency.

5. Experiments and Results

A large-scale experimental study explores the performance of the new Q^c -based AL methods. The intention is to compare those methods with each other, and to standard AL

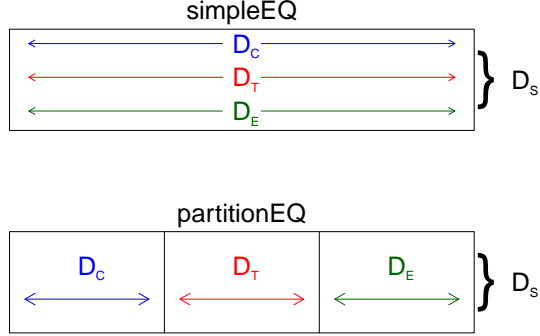


Figure 7: Different usage of the labelled data D_S by two Q^c -estimation algorithms. SimpleEQ uses all of D_S to estimate all three components; this is naïve reuse. PartitionEQ divides D_S into disjoint subsets, then uses one subset per component; this is data partitioning.

methods from the literature (described in Section 2.3). The focus is on the relative classifier improvements of each AL method, rather than absolute classifier performance.

The base classifier is varied, since AL performance is known to depend substantially on the classifier (Guyon et al., 2011; Evans et al., 2013). To provide model diversity, the study uses several classifiers with different potential capabilities: LDA, 5-nn, naïve Bayes and SVM (see Section 2.1). The classifier implementation details are described in Appendix C.

Many different classification problems are explored, including real and simulated data, described in Appendix B. These problems are divided into three problem groups to clarify the results, see Section 5.4. The experimental study uses error rate for the loss function L (see Section 2.1). Further results are available for other loss functions (Brier Score and the H-measure) but are omitted for space.

The definition of Q^c in Equation 2 conditions on certain sources of variation: the covariate location \mathbf{x} , the classifier θ , the labelled dataset D_S , and the true class probability vector \mathbf{p} . The experimental study explores several sources of variation: the AL algorithms, the classifier θ , and the classification problem (\mathbf{X}, Y) .

5.1 Active Learning Methods

The experimental study evaluates many AL methods, to compare their performance across a range of classification problems. These methods fall into three groups: RS as the natural benchmark of AL, standard AL methods from the literature, and algorithms estimating Q^c . The second group consists of four standard AL methods: SE, QbcV, QbcA, and EfeLc

(all described in Section 2.3). The third group contains the two Q^c -estimation algorithms, simpleEQ and partitionEQ, defined in Section 4 and abbreviated as SEQ and PEQ.

For the two Qbc methods, a committee of four classifiers is chosen for model diversity: logistic regression, 5-nn, 21-nn, and random forest. Logistic regression is a parametric discriminative classifier described in Schein and Ungar (2007); random forest is a non-parametric classifier described in Breiman (2001); K -nn is described in Section 2.1. This committee is arbitrary, but diverse; the choices of committee size and constitution are open research problems.

Density weighting is sometimes recommended in the AL literature, see Olsson (2009). However, the effects of density weighting are not theoretically understood. The experimental study also generated results from density weighting, omitted due to space, which left unaltered the primary conclusion that Q^c -estimation algorithms are competitive with standard methods from the literature. The issue of density weighting is deferred to future work.

5.2 Experimental AL Sandbox

Iterated AL provides for the exploration of AL performance across the whole learning curve, see Section 2.2 and Guyon et al. (2011); Evans et al. (2013). In this experimental study, the AL iteration continues until the entire pool has been labelled. Each single experiment has a specific context: a classification problem, a base classifier, and a random seed. Monte Carlo replication is applied to the experiments via the random seed, which is used to reshuffle the classification data. The seed thereby affects the pool X_P and the test data D_E , which are both drawn from the reshuffled classification data; ten seeds are used.

Given this experimental context, the experimental AL sandbox then evaluates the performance of all AL methods over a single dataset, using iterated AL. Each AL method produces a learning curve that shows the overall profile of loss as the number of labelled examples increases. To illustrate, Figure 8 shows the learning curve for several AL methods.

5.3 Assessing Performance

The AL literature provides a selection of metrics to assess AL performance, such as AUA (Guyon et al., 2011), WI (Evans et al., 2013) and label complexity (Dasgupta, 2011). Experimental results suggest substantial agreement between the metrics, in their ranking of AL methods. For that reason, and to avoid any arbitrary choice of one single metric, this study employs several metrics.

The experimental study evaluates four metrics: AUA, WI with two weighting functions (exponential with $\alpha = 0.02$, and linear), and label complexity (with $\epsilon = 5$). The *overall rank* is also calculated as the ranking of the mean ranks. This yields five AL metrics: four primary metrics (label complexity, AUA, WI-linear, WI-exponential) and one aggregate metric (overall rank).

As discussed in Section 2.2, AL performance metrics assess the relative improvements in classifier performance, when comparing one AL method against another (or when comparing AL against RS). Thus the real quantity of interest is the ranking of the AL methods.

For a single experiment, there is a single classification problem and base classifier. In such an experiment, all five metrics are evaluated for every AL method, so that each metric




Figure 8: Result for a single experiment of iterated AL. Each AL method performs multiple selection steps, generating a set of losses that define the learning curve. For clarity, a smoothed representation of the data is presented. The early part of the learning curve is shown; beyond that, all learning curves tend to the same asymptote. The classification problem is the Ripley four-Gaussian problem (see Figure 6 and Appendix B), with the base classifier being 5-nn.

produces its own ranking of the AL methods. Since there are seven AL methods (see Section 5.1), the ranks fall between one and seven, with some ties.

The results show that the AL metrics substantially agree on AL method ranking (see Tables 1 and 2). This agreement suggests that the ranking results are reasonably insensitive to the choice of AL metric. In this experimental study, AL performance is assessed by overall rank.

5.4 Assembly of Aggregate Results

To address the variability of AL, multiple Monte Carlo draws are conducted for each dataset. Thus for each experiment, the labelled, pool and test data are drawn from the population, as different independent subsamples. This random subsampling addresses two primary sources of variation, namely the initially labelled data and the unlabelled pool.

The experimental study examines many Monte Carlo draws, classification problems in groups, and classifiers. The aggregate results are calculated by averaging, first over Monte Carlo replicates, and then over classification problems.

Table 1: Results for a single pairing of classifier and problem. The base classifier is 5-nn. The classification problem is Credit-95. The six methods shown are selected as the best by overall rank.

Classifier 5-nn, and Credit-95 Problem						
AL Performance Metric	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
	SE	QbcV	SEQ	PEQ	RS	QbcA
Label Complexity	4	2	1	6	2	5
AUA	1	2	3	4	5	6
WI-Linear	1	2	6	3	5	4
WI-Exponential	1	2	4	3	6	5
Overall Rank	1	2	3	4	5	6

Table 2: Results for a single classifier and a group of problems. The base classifier is 5-nn. The classification problem group is the large problem group (see Appendix B). The six methods shown are selected as the best by overall rank.

Classifier 5-nn, and large data problem group						
AL Performance Metric	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
	SE	PEQ	SEQ	QbcV	QbcA	RS
Label Complexity	7	4	2	3	5	1
AUA	1	3	2	4	5	6
WI-Linear	1	2	5	4	3	6
WI-Exponential	1	2	3	4	5	6
Overall Rank	1	2	3	4	5	6

First the results over the Monte Carlo replicates are averaged, to produce the aggregate result for a given pairing of classifier and problem. Table 1 is shown just to illustrate one such result, where six AL methods are chosen as the best by overall rank. Second, further averaging over the problem within a group yields the aggregate result for a pairing of classifier and problem group; this second averaging is shown in Table 2. Finally the results are assembled for all three problem groups, to form the aggregate results for a single classifier, shown in Table 3. For each group of problems, six methods are judged as the best by overall rank.

Table 3: Results for base classifier LDA over three groups of problems. The six methods shown are selected as the best by overall rank.

Classifier LDA						
	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
Small Problems	QbcA	QbcV	SE	PEQ	SEQ	RS
Large Problems	PEQ	QbcV	SE	RS	SEQ	QbcA
Theoretical Problems	QbcV	SE	PEQ	SEQ	QbcA	RS

5.5 Results

Table 4: Results for base classifier 5-nn over three groups of problems. The six methods shown are selected as the best by overall rank.

Classifier 5-nn						
	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
Small Problems	SE	RS	QbcV	QbcA	PEQ	SEQ
Large Problems	SE	PEQ	SEQ	QbcA	QbcV	RS
Theoretical Problems	PEQ	QbcV	SE	RS	SEQ	QbcA

Table 5: Results for base classifier naïve Bayes over three groups of problems. The six methods shown are selected as the best by overall rank.

Classifier naïve Bayes						
	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
Small Problems	RS	QbcV	SE	SEQ	PEQ	QbcA
Large Problems	QbcV	SE	SEQ	PEQ	RS	QbcA
Theoretical Problems	QbcV	SE	QbcA	RS	PEQ	SEQ

The results for LDA, K -nn, naïve Bayes and SVM are shown in Tables 3, 4, 5 and 6 respectively. These results are the primary results of the of the experimental study, covering the four classifiers, all the problems in three groups, and multiple Monte Carlo replicates.

For all four classifiers, the Q^c -motivated algorithms perform effectively in comparison to the standard AL methods from the literature. This conclusion holds true over different

classifiers and different classification problems. This reinforces the advantages of the explicit theoretical target Q^c .

Considering the AL methods from the literature, QBC and SE consistently perform well. For QBC, vote entropy (QbcV) mostly outperforms average Kullback-Liebler divergence (QbcA). EfeLc performs somewhat less well, perhaps because of the way it approximates loss using the unlabelled pool (see Section 2.3).

Comparing the Q^c -estimation algorithms, the algorithm partitionEQ outperforms the algorithm simpleEQ, in all cases except two. The fact that the data partitioning approach of partitionEQ outperforms the naive reuse approach of simpleEQ may have implications for Q^c -estimation algorithms in general. This suggests that more sophisticated forms of data partitioning and resampling may improve Q^c -estimation algorithms. Such algorithms are the subject of future work.

6. Conclusion

Example quality is a new theoretical approach to AL which defines optimal behaviour via loss. This provides the abstract target Q^c , which is defined for any classifier and for any loss function.

Exploring this abstract definition of optimal AL behaviour generates new insights into AL. For example, the optimal selection is examined and compared to known AL methods, revealing suboptimal choices. Further, the dependence of the target Q^c on the labelled data motivates the definition and exploration of optimal behaviour independent of the labelled data.

While the approach is primarily theoretical, example quality also addresses AL applications, by defining Q^c as a target for algorithms to estimate. A comprehensive experimental study compares the performance of Q^c -estimation algorithms alongside several standard AL methods. The results demonstrate that the Q^c algorithms are competitive across a range of classifiers and problems. Thus example quality provides insights into AL, and motivates methods to address example selection in practice.

This framework lays the ground for future theoretical examination of other aspects of AL, particularly density weighting. This work defines two straightforward algorithms to estimate Q^c ; more sophisticated estimators are the subject of future research.

Acknowledgments

Table 6: Results for base classifier SVM over three groups of problems. The six methods shown are selected as the best by overall rank.

Classifier SVM						
	Method 1	Met. 2	Met. 3	Met. 4	Met. 5	Met. 6
Small Problems	QbcV	QbcA	PEQ	RS	SE	SEQ
Large Problems	QbcV	QbcA	PEQ	SE	SEQ	RS
Theoretical Problems	QbcV	RS	PEQ	QbcA	SEQ	SE

The work of Lewis P. G. Evans is supported by a doctoral training award from the EPSRC.

Appendix A.

Section 3.2 describes a theoretical classification problem and calculates Q^c as an explicit function of x for error loss L_e . This Appendix presents an analytic calculation of the target Q^m for error loss L_e , and a brief exploration of the behaviour of Q^m . Finally the calculation of error loss L_e is given. Further results are available but omitted for space, such as the targets Q^c and Q^m for other loss functions (quadratic loss and Brier Score), see Hand (1997).

Analytic Calculation of Q^m with Error Loss

This is the analytic calculation of Q^m under error loss L_e . In this Appendix, n_s is abbreviated to n .

To recap from Section 3.2, $(X|Y = c_1) \sim N(-1, 1)$, $(X|Y = c_2) \sim N(1, 1)$ and $\hat{t} \sim N(0, \frac{1}{n})$.

The true densities $q_j(x)$ and estimated densities $\hat{q}_j(x)$ of the Gaussian class-conditional densities $(X|Y = c_j)$ are given by

$$q_j(x) = \frac{1}{\sqrt{2\pi}} \exp \frac{-1}{2}(x - \mu_j)^2 \quad ; \quad \hat{q}_j(x) = \frac{1}{\sqrt{2\pi}} \exp \frac{-1}{2}(x - \hat{\mu}_j)^2.$$

The true class probabilities $p_j(x)$ and estimated class probabilities $\hat{p}_j(x)$ of the Gaussian class-conditional densities $(X|Y = c_j)$ are given by

$$p_j(x) = \frac{q_j(x)}{q_1(x) + q_2(x)} \quad ; \quad \hat{p}_j(x) = \frac{\hat{q}_j(x)}{\hat{q}_1(x) + \hat{q}_2(x)}.$$

Further, $\hat{t}'_j = \hat{t} + \frac{z}{2}(x - \hat{\mu}_j) = (\hat{t} + \frac{zx}{2}) - (\frac{z\hat{\mu}_j}{2})$. For class c_1 , $\hat{t}'_1 = (\hat{t} + \frac{zx}{2}) - (\frac{z\hat{\mu}_1}{2})$; of those final two terms, the sampling distributions are $N(\frac{zx}{2}, \frac{1}{n})$ and $N(\frac{-z}{2}, \frac{2z^2}{4n})$ respectively. Hence the sampling distribution of \hat{t}'_1 is $\hat{t}'_1 \sim N(\frac{z}{2}(x + 1), \frac{2+z^2}{2n})$.

Whereas for class c_2 , $\hat{t}'_2 = (\hat{t} + \frac{zx}{2}) - (\frac{z\hat{\mu}_2}{2})$; of those final two terms, the sampling distributions are $N(\frac{zx}{2}, \frac{1}{n})$ and $N(\frac{z}{2}, \frac{2z^2}{4n})$ respectively. Hence the sampling distribution of \hat{t}'_2 is $\hat{t}'_2 \sim N(\frac{z}{2}(x - 1), \frac{2+z^2}{2n})$.

First consider the expectation of Term F. Equation 4 gives

$$L_e(\hat{\mu}_1, \hat{\mu}_2) = \frac{1}{2} \{1 - F_1(\hat{t}) + F_2(\hat{t}) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}_2)[2F_1(\hat{t}) - 2F_2(\hat{t})]\}.$$

For F_1 , $F_1(\hat{t}) = P(q_1(x) < \hat{t}) = P(q_1(x) - \hat{t} < 0) = F_{Z_{10}}(0)$ for $Z_{10} = q_1(x) - \hat{t}$. The sampling distributions are given by $q_1 \sim N(-1, 1)$ and $\hat{t} \sim N(0, \frac{1}{n})$. Hence $Z_{10} \sim N(-1, \frac{n+1}{n})$.

For F_2 , $F_2(\hat{t}) = P(q_2(x) < \hat{t}) = P(q_2(x) - \hat{t} < 0) = F_{Z_{11}}(0)$ for $Z_{11} = q_2(x) - \hat{t}$. The sampling distributions are given by $q_2 \sim N(+1, 1)$ and $\hat{t} \sim N(0, \frac{1}{n})$. Hence $Z_{11} \sim N(+1, \frac{n+1}{n})$.

Further, $P(\mathbb{1}(\hat{\mu}_1 > \hat{\mu}_2)) = P(\hat{\mu}_1 > \hat{\mu}_2) = P(\hat{\mu}_1 - \hat{\mu}_2 > 0) = 1 - F_{Z_{12}}(0)$, for $Z_{12} = \hat{\mu}_1 - \hat{\mu}_2$, hence $Z_{12} \sim N(-2, \frac{4}{n})$. Combining these sampling distributions with the definition of Term

F as $L_e(\hat{\mu}_1, \hat{\mu}_2)$ yields

$$E[\text{Term F}] = \frac{1}{2}\{1 - F_{Z_{10}}(0) + F_{Z_{11}}(0) + [1 - F_{Z_{12}}(0)][2F_{Z_{10}}(0) - 2F_{Z_{11}}(0)]\}.$$

Turning to Term J which is given by

$$\begin{aligned} \text{Term J} = & \frac{p_1}{2}\{1 - F_1(\hat{t}'_1) + F_2(\hat{t}'_1) + \mathbb{1}(\hat{\mu}'_1 > \hat{\mu}_2)[2F_1(\hat{t}'_1) - 2F_2(\hat{t}'_1)]\} \\ & + \frac{p_2}{2}\{1 - F_1(\hat{t}'_2) + F_2(\hat{t}'_2) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}'_2)[2F_1(\hat{t}'_2) - 2F_2(\hat{t}'_2)]\}. \end{aligned}$$

The sampling distributions of \hat{t}'_1 and \hat{t}'_2 are given above. $F_1(\hat{t}'_1) = P(q_1(x) < \hat{t}'_1) = P(q_1(x) - \hat{t}'_1 < 0) = F_{Z_3}(0)$ for $Z_3 = q_1(x) - \hat{t}'_1$. Hence $Z_3 \sim N(-1 - \frac{z}{2}(x+1), 1 + \frac{2+z^2}{2n})$. $F_2(\hat{t}'_1) = P(q_2(x) < \hat{t}'_1) = P(q_2(x) - \hat{t}'_1 < 0) = F_{Z_6}(0)$ for $Z_6 = q_2(x) - \hat{t}'_1$. Hence $Z_6 \sim N(1 - \frac{z}{2}(x+1), 1 + \frac{2+z^2}{2n})$. Further, $P(\mathbb{1}(\hat{\mu}'_1 > \hat{\mu}_2) = P(\hat{\mu}'_1 - \hat{\mu}_2 > 0) = 1 - F_{Z_{13}}(0)$, for $Z_{13} = \hat{\mu}'_1 - \hat{\mu}_2$, hence $Z_{13} \sim N(z + zx - 2, \frac{2}{n}(z^2 - 2z + 2))$.

Further, $F_1(\hat{t}'_2) = P(q_1(x) < \hat{t}'_2) = P(q_1(x) - \hat{t}'_2 < 0) = F_{Z_8}(0)$ for $Z_8 = q_1(x) - \hat{t}'_2$. Hence $Z_8 \sim N(-1 - \frac{z}{2}(x-1), 1 + \frac{2+z^2}{2n})$. $F_2(\hat{t}'_2) = P(q_2(x) < \hat{t}'_2) = P(q_2(x) - \hat{t}'_2 < 0) = F_{Z_9}(0)$ for $Z_9 = q_2(x) - \hat{t}'_2$. Hence $Z_9 \sim N(1 - \frac{z}{2}(x-1), 1 + \frac{2+z^2}{2n})$. Further, $P(\mathbb{1}(\hat{\mu}_1 > \hat{\mu}'_2) = P(\hat{\mu}_1 - \hat{\mu}'_2 > 0) = 1 - F_{Z_{16}}(0)$, for $Z_{16} = \hat{\mu}_1 - \hat{\mu}'_2$, hence $Z_{16} \sim N(z - zx - 2, \frac{2}{n}(z^2 - 2z + 2))$.

Using these sampling distributions with the definition of Term J above yields

$$\begin{aligned} Q^m(x, n) = & E[\text{Term F}] - E[\text{Term J}] \\ = & \frac{1}{2}\{1 - F_{Z_{10}}(0) + F_{Z_{11}}(0) + [1 - F_{Z_{12}}(0)][2F_{Z_{10}}(0) - 2F_{Z_{11}}(0)]\} \\ & - \frac{p_1}{2}\{1 - F_{Z_3}(0) + F_{Z_6}(0) + [1 - F_{Z_{13}}(0)][2F_{Z_3}(0) - 2F_{Z_6}(0)]\} \\ & - \frac{p_2}{2}\{1 - F_{Z_8}(0) + F_{Z_9}(0) + [1 - F_{Z_{16}}(0)][2F_{Z_8}(0) - 2F_{Z_9}(0)]\}, \end{aligned}$$

where the distributions of the Z_i variables are given above.

Figure 9 illustrates Q^m as a function of x . This expression for Q^m is a complicated non-linear function of x , due to the non-linearity of $p_1(x)$, $p_2(x)$ and $F_{Z_i}(0)$. The substantial set of central x values where Q^m is positive shows that many choices of x will improve the classifier.

Analytic Calculation of Error Loss

First the classifier decision rule $r_1(x)$ given in Section 3.2.1 can be expressed as

$$r_2(x) = \begin{cases} \hat{y} = c_1 & : x < \hat{t}, \\ \hat{y} = c_2 & : x > \hat{t}, \end{cases} \quad r_3(x) = \begin{cases} \hat{y} = c_1 & : x > \hat{t}, \\ \hat{y} = c_2 & : x < \hat{t}, \end{cases} \quad r_1(x) = \begin{cases} \hat{y} = r_2(x) & : \hat{\mu}_1 < \hat{\mu}_2, \\ \hat{y} = r_3(x) & : \hat{\mu}_1 > \hat{\mu}_2, \end{cases}$$

The probability of an error for a single covariate x is denoted $p(e|x) = p(y = c_1, \hat{y} = c_2|x) + p(y = c_2, \hat{y} = c_1|x)$. Given x , the probabilities of y and \hat{y} are independent. Thus for decision rule r_2 , the loss L_e is given by




Figure 9: Illustration of the target Q^m as a function of x . The green line indicates $Q^m(x) = 0$ (zero improvement), with n_s being 18. The substantial set of central x values where Q^m is positive shows that many choices of the example x will improve the classifier.

$$\begin{aligned}
 L_e(\hat{\theta}) &= E_X[p(e|x)] = \int p(y = c_1, \hat{y} = c_2|x)p(x)dx + \int p(y = c_2, \hat{y} = c_1|x)p(x)dx \\
 &= \int p_1(x)1(x > \hat{t})p(x)dx + \int p_2(x)1(x < \hat{t})p(x)dx \\
 &= \int \frac{1}{2}q_1(x)1(x > \hat{t})dx + \int \frac{1}{2}q_2(x)1(x < \hat{t})dx = \frac{1}{2}[\int_{\hat{t}}^{\infty} q_1(x)dx + \int_{-\infty}^{\hat{t}} q_2(x)dx] \\
 &= \frac{1}{2}[1 - F_1(\hat{t}) + F_2(\hat{t})].
 \end{aligned}$$

Whereas for decision rule r_3 , the loss L_e is given by $L_e(\hat{\theta}) = \frac{1}{2}[1 + F_1(\hat{t}) - F_2(\hat{t})]$. Hence the loss L_e in general is given by

$$L_e(\hat{\theta}) = \frac{1}{2}\{1 - F_1(\hat{t}) + F_2(\hat{t}) + \mathbb{1}(\hat{\mu}_1 > \hat{\mu}_2)[2F_1(\hat{t}) - 2F_2(\hat{t})]\}.$$

Appendix B.

A diverse set of classification problems is chosen to explore AL performance. The classification problems fall into two sets: real problems and theoretical problems.

First the real data classification problems are shown in Tables 7 and 8. The real data problems are split into two groups, one for smaller problems of fewer examples, and another of larger problems. The class prior is shown, since the experimental study uses error rate as loss. The sources for this data include UCI (Bache and Lichman, 2013), Guyon et al. (2011), Anagnostopoulos et al. (2012) and Adams et al. (2010).

The intention here is to provide a wide variety in terms of problem properties: covariate dimension d , number of classes k , the class prior π , and the underlying distribution. The number and variety of problems suggests that the results in Section 5 have low sensitivity to the presence or absence of one or two specific problems.

Table 7: Real Data Classification Problems, Smaller

Name	Dim. d	Classes k	Cases n	Class Prior π
Abalone	8	3	4177	(0.31, 0.32, 0.37)
Australian	14	2	690	(0.44, 0.56)
Balance	4	3	625	(0.08, 0.46, 0.46)
Breast Cancer Wdbc	4	3	569	(0.37, 0.63)
German Credit	24	2	1000	(0.3, 0.7)
Glass	10	6	214	(0.33,0.36,0.08,0.06,0.04,0.14)
Heart-Statlog	13	2	270	(0.65, 0.44)
Ionosphere	34	2	351	(0.64, 0.36)
Monks-1	6	2	432	(0.5, 0.5)
Monks-2	6	2	432	(0.5, 0.5)
Monks-3	6	2	432	(0.5, 0.5)
Pima Diabetes	8	2	768	(0.35, 0.65)
Sonar	60	2	208	(0.47, 0.53)
Wine	13	3	178	(0.33, 0.4, 0.27)

Table 8: Real Data Classification Problems, Larger

Name	Dim. d	Classes k	Cases n	Class Prior π
Fraud	20	2	5999	(0.167, 0.833)
Electricity Prices	6	2	27552	(0.585, 0.415)
Colon	16	2	17076	(0.406, 0.594)
Credit 93	29	2	4406	(0.007, 0.993)
Credit 94	29	2	8493	(0.091, 0.909)
Credit 95	29	2	21076	(0.117, 0.883)
Credit 96	29	2	24396	(0.111, 0.889)
Credit 97	29	2	21191	(0.043, 0.957)
AL Challenge, Alex	11	2	10000	(0.270, 0.730)
AL Challenge, Sina	92	2	20722	(0.062, 0.938)

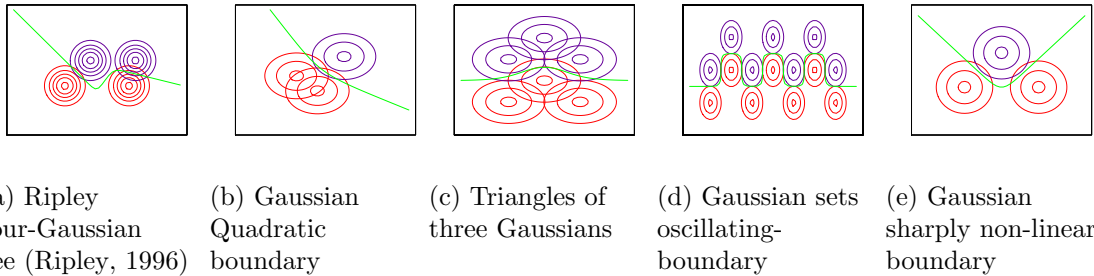


Figure 10: Density contour plots showing the theoretical classification problems. The class-conditional distributions $(\mathbf{X}|y = c_j)_1^k$ are shown in red for class 1 and blue for class 2. These class-conditional distributions $(\mathbf{X}|y = c_j)_1^k$ are either Gaussians or mixtures of Gaussians. The decision boundary is shown in green.

Second the theoretical classification problems are illustrated in Figure 10. These theoretical problems are generated by sampling from known probability distributions. The class-conditional distributions $(\mathbf{X}|y = c_j)_1^k$ are either Gaussians or mixtures of Gaussians. This set of problems presents a variety of decision boundaries to the classifier. All have balanced uniform priors, and the Bayes Error Rates are approximately 0.1.

Appendix C.

Here the classifier implementation details are described. For LDA, the standard R implementation is used. For 5-nn, the R implementation from package `knn` is used. This implementation applies covariate scaling: each covariate is scaled to have equal standard deviation (using the same scaling for both training and testing data). For naïve Bayes, the R implementation from package `e1071` is used. For metric predictors, a Gaussian distribution (given the target class) is assumed. This approach is less than ideal, but tangential to the theoretical target and experimental study. For SVM, the R implementation from package `e1071` is used. The SVM kernel used is radial basis kernel. The probability calibration of the scores is performed for binary problems by MLE fitting of a logistic distribution to the decision values, or for multi-class problems, by computing the a-posteriori class probabilities using quadratic optimisation.

References

Niall M. Adams, Dimitris K. Tasoulis, Christoforos Anagnostopoulos, and David J. Hand. Temporally-adaptive linear classification for handling population drift in credit scoring. In *Proceedings of the 19th International Conference on Computational Statistics*, pages 167–176, 2010.

Christoforos Anagnostopoulos, Dimitrios Tasoulis, Niall M. Adams, and David J. Hand. Online linear and quadratic discriminant analysis with adaptive forgetting for streaming classification. *Statistical Analysis and Data Mining*, 5:139–166, 2012.

- Kevin Bache and Moshe Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- Sanjoy Dasgupta. Two faces of active learning. *Theoretical Computer Science*, 412(19):1767–1781, 2011.
- Sanjoy Dasgupta and Daniel Hsu. Hierarchical sampling for active learning. In *Proceedings of the 25th International Conference of Machine Learning*, pages 208–215, 2008.
- Richard Duda, Peter Hart, and David Stork. *Pattern Classification*. Wiley, 2001.
- Lewis P. G. Evans, Niall M. Adams, and Christoforos Anagnostopoulos. When does active learning work? In *Intelligent Data Analysis Conference Proceedings, Lecture Notes in Computer Science Series by Springer-Verlag*, pages 174–185, 2013.
- Robert C. Geary. The contiguity ratio and statistical mapping. *The Incorporated Statistician*, 5(3):115–127, 129–146, 1954.
- Baohua Gu, Feifang Hu, and Huan Liu. Modelling classification performance for large data sets. *Advances in Web-Age Information Management, Lecture Notes in Computer Science*, 2118:317–328, 2001.
- Isabelle Guyon, Gavin Cawley, Gideon Dror, and Vincent Lemaire. Results of the active learning challenge. *Journal of Machine Learning Research*, 16:19–45, 2011.
- David Hand. *Construction and Assessment of Classification Rules*. Wiley, 1997.
- David Hand and Keming Yu. Idiot’s Bayes: Not so stupid after all? *International Statistical Review*, 69(3):385–398, 2001.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, 2nd edition, 2009.
- Steven C. H. Hoi, Rong Jin, Jianke Zhu, and Michael R. Lyu. Batch mode active learning and its application to medical image classification. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 417–424, 2006.
- Carl Myers Kadie. Seer: Maximum likelihood regression for learning-speed curves. Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Computer Science in the Graduate College of the University of Illinois at Urbana-Champaign, 1995.
- Mohit Kumar, Rayid Ghani, Mohak Shak, Jaime Carbonell, and Alexander Rudnick. Empirical comparison of active learning strategies for handling temporal drift. *ACM Transactions on Embedded Computing Systems*, 9(4):161–168, 2010.

- Tom Mitchell. *Machine Learning*. McGraw-Hill Higher Education, 1997.
- Patrick A. P. Moran. Notes on continuous stochastic phenomena. *Biometrika*, 37:17–23, 1950.
- Fredrik Olsson. A literature survey of active machine learning in the context of natural language processing. Technical Report ISSN: 1100-3154, Swedish Institute of Computer Science, 2009.
- Claudia Perlich, Foster Provost, and Jeffrey S. Simonoff. Tree induction vs. logistic regression: A learning-curve analysis. *Journal of Machine Learning Research*, 4:211–255, 2003.
- Foster Provost and Josh Attenberg. Inactive learning difficulties employing active learning in practice. *ACM SIGKDD*, 12:36–41, 2010.
- Brian D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, 1996.
- Nicholas Roy and Andrew McCallum. Toward optimal active learning through sampling estimation of error reduction. In *Proceedings of the 18th International Conference on Machine Learning*, pages 441–448, 2001.
- Andrew Schein and Lyle Ungar. Active learning for logistic regression: an evaluation. *Machine Learning*, 68(3):235–265, 2007.
- Burr Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
- H. Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Proceedings of the 5th annual workshop on Computational Learning Theory*, pages 287–294, 1992.
- Sebastian Thrun and Knut Möller. Active exploration in dynamic environments. In *Proceeding of the 25th Conference of Advances in Neural Information Processing Systems*, volume 4, pages 531–538, 1992.